

CLAIMS

What is claimed is:

1. A method for use in a system for storing and accessing data, the system including at least one initiator, at least one target, and at least one switch, the method comprising:
- providing a plurality of paths to the target from the initiator, each path passing through the switch;
 - 10 dynamically load balancing amongst the paths by the switch.
2. The method of claim 1, wherein dynamically load balancing amongst the paths includes:
- determining a respective average response time for each path;
 - 15 passing a request received by the switch from the initiator to the target along the path with the shortest average response time.
3. The method of claim 1, wherein the target is a physical storage device.
- 20 4. The method of claim 1, wherein the target is a virtual target.
5. The method of claim 1, wherein the target is a mirrored target with a plurality of members and wherein load balancing amongst the paths includes:
- determining a respective average response time of each member of the
 - 25 mirrored target;
 - passing a request received by the switch from the initiator to the target to the member with the shortest average response time.
6. The method of claim 5, wherein the request is a read request.

30

7. The method of 1, wherein the switch includes a plurality of ports and wherein load balancing is performed by respective circuitry affiliated with each respective port.

5 8. A method for use in a storage network including an initiator, a storage device, and a switch, the method comprising:

providing a plurality of paths from the storage device to the initiator, each path passing through the switch;

determining a respective average response time for each path;

10 passing a request received by the switch from the initiator to the storage device along the path with the shortest average response time.

9. A method for use in a storage network including an initiator, a mirrored virtual target having a plurality of members, and a switch, the method comprising:

15 providing a path from each member of the mirrored virtual target to the initiator, each path passing through the switch;

determining a respective average response time for each path;

20 passing a request received by the switch from the initiator to the member with the shortest average response time.

10. The method of claim 9, wherein the request is a read request.

25 11. A method for use in a storage network including a switch, a plurality of initiators, and a plurality of targets, wherein some of the targets are mirrored targets with a plurality of members and some of the targets are physical storage devices, the method comprising:

providing a plurality of paths from a first initiator to a physical storage device via the switch;

providing a respective path from a second initiator to each member of a mirrored target via the switch;

determining a respective average response time for each path from the first initiator to the physical storage device and for each path from the second initiator to each member of the mirrored target;

passing a first request received by the switch from the initiator to the physical storage device along the path to the physical storage device with the shortest average response time;

passing a second request received by the switch from the initiator to the member of the mirrored target with the shortest average response time.

12. The method of claim 11, wherein:

the step of passing a first request is performed by a first linecard in the switch; and

the step of passing a second request is performed by a second linecard in the switch.

13. The method of claim 11, wherein the step of passing a first request and the step of passing a second request are both performed by the same linecard.

14. The method of claim 11, wherein the switch includes a plurality of linecards and wherein the step of determining is performed by each linecard.

15. A switch for use in a storage network, comprising:
a plurality of ports,
load balancing circuitry affiliated with each of the ports.

16. The switch of claim 15, wherein:
the load balancing circuitry includes a storage processor and a CPU.

17. A switch for use in a storage network, the network including an initiator and a target in communication with the initiator by a plurality of paths, each path passing through the switch, the switch comprising:

a plurality of ports;

means for load balancing amongst the paths.

18. The switch of claim 17, wherein the means for load balancing includes:

means for maintaining statistics for the response time of each path;

means for passing a request received by the switch from the initiator to the target along the path with the shortest average response time.

19. A storage network, including:

an initiator;

a target;

a switch;

a plurality of paths from the initiator to the target via the switch;

wherein the switch includes statistical information regarding the response time for each path; and

wherein the switch is designed to forward a request from the initiator to the target along the path with the shortest response time.

20. The storage network of claim 19, wherein the target is a physical storage device.

21. The storage network of claim 19, wherein the target is a virtual target.

22. The storage network of claim 19, wherein the target is a mirrored target with a plurality of members and wherein the plurality of paths are respective paths to each member.

23. A machine readable media which has instructions stored thereon, which when executed by a switch in a storage network causes the switch to perform the following steps:

providing a plurality of paths to a target from an initiator, each path
5 passing through the switch;
determining a respective response time of each path;
passing a request received by the switch from the initiator to the target
along the path with the shortest average response time.

24. The machine readable media of claim 23, wherein the target is a physical
10 storage device.

25. The machine readable media of claim 23, wherein the target is a virtual
15 target.

26. The machine readable media of claim 23, wherein the target is a mirrored
target with a plurality of members and wherein the instructions to further
include:

determining a respective response time of each member of the mirrored
20 target;
passing a request received by the switch from the initiator to the target to
the member with the shortest average response time.

20051053.011802

ai